Boosting Neural Machine Translation

Dakun Zhang and **Jungi Kim** and **Josep Crego** and **Jean Senellart**
firstname.lastname@systrangroup.com
SYSTRAN / 5 rue Feydeau, 75002 Paris, France

arXiv:1612.06138v1 [cs.CL] 19 Dec 2016

Abstract

Training efficiency is one of the main problems for Neural Machine Translation (NMT). Deep networks, very large data and many training iterations are necessary to achieve state-of-the-art performance for NMT. This results in very high computation cost and slow down research and industrialization. In this paper, we first investigate the instability by randomizations for NMT training, and further propose an efficient training method based on data boosting and bootstrapping with no modifications to the neural network. Experiments show that this method can converge much faster compared with a baseline system and achieve stable improvement up to 2.36 BLEU points with 80% training cost.

1 Introduction

With the rapid development of research on Neural Machine Translation (NMT), translation quality has been improved significantly compared with traditional statistical based method (Bahdanau et al., 2014; Cho et al., 2015; Zhou et al., 2016; Sennrich et al., 2015). However, the training efficiency is one of the main problems for both academia and industry. It may take several weeks to learn a well-trained NMT model with a medium size corporus even with cutting-edge hardware.

NMT system directly models the mappings between source sentence $x_1^n = (x_1, ..., x_n)$ and target sentence $y_1^m = (y_1, ..., y_m)$, with n and m words respectively (Sutskever et al., 2014). NMT system is usually based on an encoder-decoder-attention framework, in which the source sentence is fed into an encoder word by word to form a fixed length representation vector, with a forward sequence of hidden states $(\overrightarrow{h_1}, ..., \overrightarrow{h_n})$ and a backward sequence $(\overleftarrow{h_1}, ..., \overleftarrow{h_n})$. With the attention mechanism (Bahdanau et al., 2014), a decoder is used to decide which part of the source sentence to pay attention to and predict corresponding word representation at time t together with history predictions before time t. Then, a softmax is used to restore the word representation to natural target words. During the training process, the parameter θ is optimized:

$$p(y_1^m|x_1^n;\theta) = \prod_{t=1}^{m} p(y_t|y_{<t}, x_1^n;\theta)$$

In order to achieve an acceptable performance on systems, deep networks (up to 8 layers) and more iterations (10-18 epochs) are necessary with a certain amount data (Wu et al., 2016; Crego et al., 2016). Since it takes several weeks to generate results, it is difficult to experiment with several different parameters, hence slowing down innovation.

While Wu et al. (2016) proposes a brute-force approach with massive data and model parallelism as a way to accelerate training, in this paper, we investigate the influence of randomization of meta-parameters for NMT model convergence. We train several diverse models by randomizing both training data (with a random seed) and learning rates. Our evaluations show that up to 2.28 BLEU points difference can be achieved through only randomization. It also shows that training instability is a severe problem for NMT.

To overcome this problem, we focus the training on the most difficult parts of the corpus, a technique borrowed from data boosting and bootstrapping. Every several epochs, we reselect 80% of the corpus with the highest perplexity (ppl.) to use for training. We achieve a maximum 2.36 BLEU points improvement (epoch10) while saving 20% training cost at the same time, which is quite a stable improvement compared with our baseline sys-

tem on French-English translations. By increasing the size of the training corpus by an additional 10%, we achieve another 0.5 BLEU points improvement when the system converges at epoch18.

2 Related work

Training efficiency has been a main concern by many researchers in the field of NMT. Data parallelism and model parallelism are two commonly used techniques to improve the speed of training (Wu et al., 2016). As a result, multiple GPUs or TPUs[1] are needed which requires additional replication and combination costs. Parallelism does not reduce the actual cost of computation. It only saves time by using more hardware.

Chen et al. (2016b) proposed a modified RNN structure with a full output layer to facilitate the training when using a large amount of data. Kalchbrenner et al. (2016) proposed ByteNet with two networks to encode the source and decode the target at the same time. To compare, our method does not modify the network itself and can be used in any NMT framework.

Other methods focus on how to reduce the parameters trained by the model (See et al., 2016). They show that with a pruning technique, 40-60% of the parameters can be pruned out. Similar methods are proposed to reduce the hidden units with knowledge distillation (teacher/student infrastructure) (Crego et al., 2016; Kim and Rush, 2016). They retrain a smaller student model using text translated from teacher models. They report a 70% reduction on the number of parameters and a 30% increase in decoding speed. Hubara et al. (2016) proposed to reduce the precision of model parameters during training and network activations, which can also bring benefits to the training efficiency.

Some related works use data manipulations with improvement in performance, but not training. Crego et al. (2016) propose to train a corpus with 5M sentences for 10 epochs and another homogeneous 5M corpus for another 10 epochs, which can lead to better accuracy than directly training a 5M corpus with 20 epochs. Our method improves the accuracy while simultaneously reducing the cost of training.

[1]Google's Tensor Processing Unit (Wu et al., 2016).

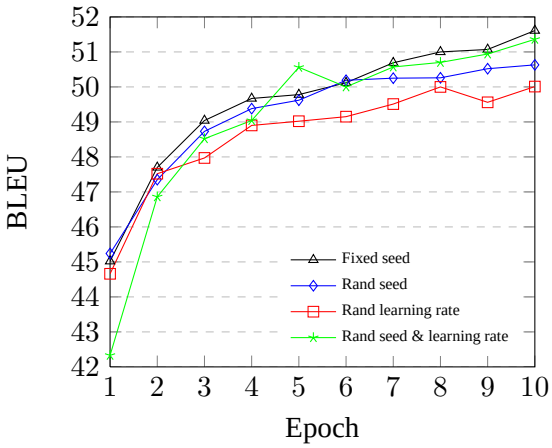

Figure 1: Effect of randomization during training for French-English translation.

3 Experiments

3.1 System Description

We build our NMT system based on `seq2seq-attn`[2], an open-source project developed by the Harvard NLP group. We use a bidirectional RNN with 4 LSTM layers with each hidden layer having 1000 nodes. We use a word embedding size of 500 and set the dropout to 0.3. The batch size is set to 64 and the default learning rate is 1.0. The maximum length of both source and target sentences is set to 80 and we limit the vocabulary size to 50K for both source and target.

3.2 Effect of Random Initialization

Randomization is one of the techniques used in training neural network models to help convergence, and to some extent, to avoid overfitting. However, there has not been much work that directly investigates the influence of randomization, especially in the field of NMT. In this experiment, we show that randomization really influences the performance of the trained model. Up to 2.28 BLEU points differences are observed.

We intentionally control the randomization of both the order of the sentences used for training and the learning rate used to control the convergence speed for the gradient descent algorithm. The system is trained on a 5M general domain French-English corpus. Each model is trained for 10 epochs and the evaluation is done on a 2,000-sentence test set with a case sensitive BLEU score. Translation is performed from French to English.

[2]`https://github.com/harvardnlp/seq2seq-attn`

There are four systems customized with:

- fixed initialization seed (default value)
- random seed for reordering sentences for each training epoch [0-10000]
- random learning rate for each training epoch [0-1][3]
- both random seed and learning rate for each training epoch

Figure 1 shows the effects of randomization. After 10 epochs when results from the different experiments tend to stabilizee, single models show a significant difference in score. It can produce a significant lower score if a bad seed is chosen. Besides that, we can see that randomizing the learning rate (red line) really influences much compared to others, and ranks the lowest for most epochs. When randomizing both the order of the sentences and the learning rate, the performance improves (green line), but it is not stable. For example, in epoch5 it gets a large boost in score, but then goes back down for epoch6. Neither randomizing just the sentence order (blue line) nor randomizing just the random learning rate (red line) has such effect.

3.3 Boosting Training

Inspired by machine learning algorithms (e.g. boosting, bootstrap, etc.), which are widely used to improve the stability and accuracy especially in the field of text classification, we force the neural network to pay closer attention to "harder" training examples. In the circumstance of NMT, we give it sentences that the system cannot translate well, which we approximate by choosing sentences with a high perplexity score. The perplexity has already been calculated and thus brings no extra calculation cost to the training.

To our knowledge, this is the first approach that integrates such meta-algorithms within NMT training. Not surprisingly, we show a steady and increased performance. To put it further, we "focused" on difficult sentences by removing easy sentences instead of "emphasizing" them by repeating them. We found improvements in both accuracy and training efficiency.

[3]Default learning rate is 1.0. We randomize this to achieve a comparable result with the 4th system with both random seed and learning rate.

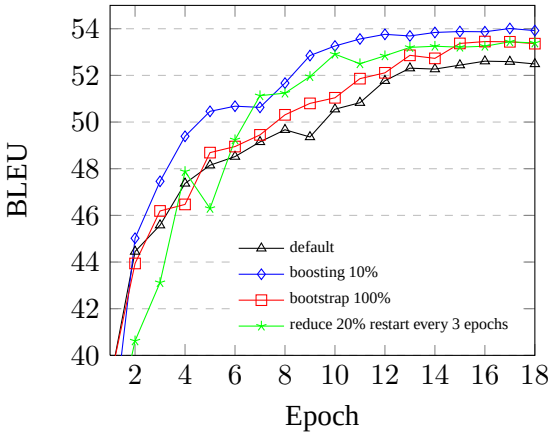

Figure 2: Effect of data boosting during training for English-French translation.

3.3.1 Single Method Boosting

In this experiment, we use a sub-set of the above mentioned corpus (1M French-English general domain) to evaluate translation in the other direction (i.e. from English to French). A detailed breakdown of the data we used are shown in Table 1. We train four systems with:

- default configuration detailed in Section 3.1
- boost 10% of the most difficult training instances ordered by ppl. for each epoch
- bootstrap (with replacement) 100% as a new training set from the base 1M corpus for each epoch
- reduce to 80% most difficult training instances ordered by ppl. for each epoch and restart this process every 3 epochs

Type	*Corpus*	*# lines*	*# src tok (EN)*	*# tgt tok (FR)*
Train	Generic	1M	24M	26M
Valid	Generic	2,000	48K	55K
Test	Generic	2,000	48K	54K

Table 1: Data used to train the model.

The system is trained up to 18 epochs[4] and the evaluation results are shown in Figure 2.

By boosting 10% of the most difficult data, the system achieved a high of **54.01** (Δ +1.52) (BLEU score), compared with the baseline system 52.49. The improvement was always stable. The training cost increased by only 10%. Comparatively, by reducing 20% of the training data, we got an

[4]For some experiments, we continue to train until 22 epochs. However, there is no further improvement.

acceptable improvement of 53.39 (Δ +0.9), with a lower training cost by about 80%. For the bootstrap method, data is selected with replacement. The score steadily improves, though not so significantly, with a final convergence of 53.36.

The best improvement happens around epoch9 and epoch10 with a score of 53.26 (Δ +2.71) compared to the baseline 50.55 for 10% extra training and by 52.91 (Δ +2.36) for using only 80% training cost. When ensembling the 4 best systems generated by single models, we get another 1.02 BLEU point improvement.

For the reduction method, every third epoch is trained using all the training data. That is to say, we use 100%, 80%, 64%, 100%, 80% ... of the training data selected by the translation difficulties (ppl.). An interesting decrease happens from epoch4 to epoch5, when the training data is reduced from 100% to 80% for a second time. A similar decrease, though not so obvious, exists from epoch7 to epoch8, epoch10 to epoch11, etc., which shows the impact of data selection.

3.3.2 Method Combination

This next step we tried to combine the methods in order to take advantage of both the accuracy and stability of the boosting and bootstrapping techniques as well as the low training costs of the reduction method. In doing so, we observed that the reduction method has the lowest cost and converges well after 3-6 epochs. For the first 3 epochs, we used boosting and bootstrap to create a relatively reasonable and stable foundation, which is important for perplexity estimation. Starting with epoch 4, we begin learning using the reduction method so as to balance both high efficiency and accuracy. However, current experiments show no improvement compared with the single method.

3.4 Averaging parameters

Junczys-Dowmunt et al. (2016) described a method to average the model parameters. They report some improvement in the experiments, though in theory, neural network models are non-linear. In our experiments, we tried to average the best models trained in Section 3.2, but we did not notice any improvement. Instead, we tried averaging the last two epochs generated from the same originations (i.e. epoch9 and epoch10 in each training configuration). An average Δ +0.31 can be achieved as in Table 2. It shows that averaging model parameters can lead to instable results.

	Epoch9	**Epoch10**	**Avg.**
Fixed seed	51.07	51.61	51.93 (+0.32)
Rand seed	50.52	50.63	51.19 (+0.56)
Rand learning rate	49.56	50.01	50.10 (+0.09)
Rand seed & learning rate	50.94	51.36	51.63 (+0.27)

Table 2: Effect of averaging models.

Though the method for averaging models is not theoretically guaranteed, it still has the advantage that it works in an offline style and brings no extra cost both for training and decoding. Therefore, we performed some other experiments by iterative training and averaging. We train the system normally for two epochs, then we average them, then train for another two epochs based on the averaged one, then average, and so on. However there is no improvement since the NMT model will re-learn the parameters based on the averaged ones. The final trained model is weaker than the model trained directly, even weaker compared to the model trained with same iterations without averaging (e.g. to compare epoch18 generated by iterative averaging with epoch12 by direct training).

4 Conclusions

For NMT, training cost is a big problem for even a medium-sized corpus with cutting-edge hardware. At the same time, the trained model is apt to converge to a local optima, which makes the training more instable. In this paper, we investigated the instability through randomizations and further proposed a data boosting method for NMT to avoid this problem. Experiments show that the improvement is quite stable during almost all training iterations. By adding 10% training corpus, translation score is improved by Δ +2.71 BLEU scores and by reducing the size of the corpus by 20%, translation performance still improved by Δ +2.36.

The method we proposed focuses on training process only. There is no restriction for the neural network structure. It can be used in any data parallelism framework and then distributed onto multi-GPUs. Also, corpus preprocessing like tokenization (e.g. using sub-word unit (Sennrich et al., 2015)) and other techniques like guided alignment (Chen et al., 2016a) can be freely added based on the method we proposed. In the future, we plan to add weights to the training instances according to their translation difficulty

to better emphasize “harder” sentences.

References

[Bahdanau et al.2014] Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. 2014. Neural machine translation by jointly learning to align and translate. *arXiv preprint arXiv:1409.0473*.

[Chen et al.2016a] Wenhu Chen, Evgeny Matusov, Shahram Khadivi, and Jan-Thorsten Peter. 2016a. Guided alignment training for topic-aware neural machine translation. *arXiv preprint arXiv:1607.01628*.

[Chen et al.2016b] Xie Chen, Xunying Liu, Yongqiang Wang, Mark JF Gales, and Philip C Woodland. 2016b. Efficient training and evaluation of recurrent neural network language models for automatic speech recognition. *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, 24(11):2146–2157.

[Cho et al.2015] Sébastien Jean Kyunghyun Cho, Roland Memisevic, and Yoshua Bengio. 2015. On using very large target vocabulary for neural machine translation.

[Crego et al.2016] Josep Maria Crego, Jungi Kim, Guillaume Klein, Anabel Rebollo, Kathy Yang, Jean Senellart, Egor Akhanov, Patrice Brunelle, Aurelien Coquard, Yongchao Deng, Satoshi Enoue, Chiyo Geiss, Joshua Johanson, Ardas Khalsa, Raoum Khiari, Byeongil Ko, Catherine Kobus, Jean Lorieux, Leidiana Martins, Dang-Chuan Nguyen, Alexandra Priori, Thomas Riccardi, Natalia Segal, Christophe Servan, Cyril Tiquet, Bo Wang, Jin Yang, Dakun Zhang, Jing Zhou, and Peter Zoldan. 2016. Systran’s pure neural machine translation systems. *CoRR*, abs/1610.05540.

[Hubara et al.2016] Itay Hubara, Matthieu Courbariaux, Daniel Soudry, Ran El-Yaniv, and Yoshua Bengio. 2016. Quantized neural networks: Training neural networks with low precision weights and activations. *arXiv preprint arXiv:1609.07061*.

[Junczys-Dowmunt et al.2016] Marcin Junczys-Dowmunt, Tomasz Dwojak, and Rico Sennrich. 2016. The amu-uedin submission to the wmt16 news translation task: Attention-based nmt models as feature functions in phrase-based smt. *arXiv preprint arXiv:1605.04809*.

[Kalchbrenner et al.2016] Nal Kalchbrenner, Lasse Espeholt, Karen Simonyan, Aaron van den Oord, Alex Graves, and Koray Kavukcuoglu. 2016. Neural machine translation in linear time. *arXiv preprint arXiv:1610.10099*.

[Kim and Rush2016] Yoon Kim and Alexander M Rush. 2016. Sequence-level knowledge distillation. *arXiv preprint arXiv:1606.07947*.

[See et al.2016] Abigail See, Minh-Thang Luong, and Christopher D Manning. 2016. Compression of neural machine translation models via pruning. *arXiv preprint arXiv:1606.09274*.

[Sennrich et al.2015] Rico Sennrich, Barry Haddow, and Alexandra Birch. 2015. Neural machine translation of rare words with subword units. *arXiv preprint arXiv:1508.07909*.

[Sutskever et al.2014] Ilya Sutskever, Oriol Vinyals, and Quoc V Le. 2014. Sequence to sequence learning with neural networks. In *Advances in neural information processing systems*, pages 3104–3112.

[Wu et al.2016] Yonghui Wu, Mike Schuster, Zhifeng Chen, Quoc V Le, Mohammad Norouzi, Wolfgang Macherey, Maxim Krikun, Yuan Cao, Qin Gao, Klaus Macherey, et al. 2016. Google’s neural machine translation system: Bridging the gap between human and machine translation. *arXiv preprint arXiv:1609.08144*.

[Zhou et al.2016] Jie Zhou, Ying Cao, Xuguang Wang, Peng Li, and Wei Xu. 2016. Deep recurrent models with fast-forward connections for neural machine translatin. *arXiv preprint arXiv:1606.04199*.